

HP OpenView

Storage Mirroring application notes

High availability for Oracle 10g

Legal and notice information

© Copyright 2005-2006 Hewlett-Packard Development Company, L.P.

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information is provided "as is" without warranty of any kind and is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Oracle® is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

Storage Mirroring High availability for Oracle 10g application notes

Document overview

This document is a Storage Mirroring application note. An application note provides guidelines on the use of Storage Mirroring in a specific environment.

This document contains:

- **Document overview**—Explains what an application note contains, how it should be used, what you need to know before trying to use the application note, and where you can go for more information.
- **Solution overview**—Explains how the applications work with Storage Mirroring and describes the considerations that you must weigh when implementing your Storage Mirroring solution. Review this section to make sure that you understand the theory involved with using Storage Mirroring and your application. Includes both basics, such as system requirements, as well as configuration and environment-specific topics, such as interactions with specific clients or special considerations for WAN (Wide Area Network) environments. Pay special attention to those topics that are directly related to your environment.
- **Sample implementations**—Describes specific examples of how to use Storage Mirroring for this solution. Use these procedures as a guideline for creating your own implementation. Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented here in order to make the solution work in your environment.

Audience

This document is written for network and application administrators who have a working understanding of the applications and environments where the Storage Mirroring solution is to be deployed. You may need to expand on the documented information in order to customize the solution to fit your environment.

Before you use this application note, you should have an understanding of:

- Storage Mirroring
- Oracle

Expectations

Application notes are intended to provide a framework for configuring a Storage Mirroring solution in a specific environment and to draw attention to decisions you will need to make when configuring your solution.

Because there are an infinite number of possible configuration, network, and environment scenarios, application notes contain general configuration guide.

This document assumes that you are comfortable working with your operating system, Storage Mirroring, and the application(s) being used with Storage Mirroring.

Related documentation

Before you begin to configure your solution, make sure that you have complete documentation for your operating system, application(s), and Storage Mirroring. This application note does not provide step-by-step instructions for using standard operating system, application, and Storage Mirroring functionality.

The following documents contain additional information that you may need while setting up this solution:

- *HP OpenView Storage Mirroring user guide* or online documentation
- Reference guides or documentation for Oracle

Getting help

Hewlett-Packard has application notes that describe how to configure Storage Mirroring with a variety of popular third-party applications. These application notes are available on the Storage Mirroring web site: <http://h18006.www1.hp.com/products/storage/software/sm/index.html>.

For help using Storage Mirroring, refer to the Storage Mirroring online manual or online help.

Solution overview

Oracle is a scalable, reliable, flexible, and high-performance relational database management system for many server-based operating systems. Storage Mirroring provides real-time enterprise data protection and replication. Storage Mirroring can be used to provide high availability for Oracle.

This document describes the steps necessary to configure Storage Mirroring to provide high availability for Windows® servers running Oracle 10g. These procedures allow a secondary server to assume the identity and role of a failed primary Oracle server while maintaining the availability of Oracle services with minimal disruption or data loss.

To complete these instructions, you will install Oracle and Storage Mirroring, and configure Storage Mirroring for replication and failover. Due to the complexities of these applications, this document is intended for network administrators with experience installing, configuring, and maintaining network applications including Storage Mirroring and Oracle.



NOTE: Storage Mirroring allows you to configure one target to monitor and failover for one or more source machines. In a one-to-one configuration, you will want to replicate your Oracle data to the same location on the target so that failover is automatic. In a many-to-one configuration, the data will need to be replicated to a unique location and then renamed to the corresponding Oracle directory on the source before failover occurs.

This application note focuses on a single Oracle server being replicated to a single target.

Requirements

- Two servers that meet one of the following operating system requirements:
 - Microsoft Windows NT 4.0 with Service Pack 4 or higher
 - Microsoft Windows 2000 with Service Pack 1 or higher
 - Microsoft Windows 2003



NOTE: The two servers should both be running the same operating system.

Modifying the sample script files

After you modify the sample scripts, save them with a new name to remove the `SAMPLE_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

Sample implementation

This section described an example of how to configure Storage Mirroring and Oracle. Use these procedures as a guideline for creating your own implementation.

Because no two environments or configurations are exactly the same, you will probably need to implement additional or different steps than what is documented in order to make the solution work in your environment.

- It is recommended that both source and target servers be member servers. (You may experience problems with promotion and demotion during failover if either of the machines are Primary or Backup Domain Controllers.)
- Both servers must be connected to the same physical network
- Two licensed copies of Oracle 10g
- Two licensed copies of Storage Mirroring

Install software on the source

1. Install Oracle on the source, if it is not already installed.
2. Install Storage Mirroring on the source using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.

Install and configure software on the target

1. Install Oracle on the target using the same installation options used on the source machine.
2. Set the Oracle service to manual startup so that the failover and failback scripts that you will be creating will control the stopping and starting of the Oracle service.
 - a. In **Control Panel, Administrative Tools, Services**, double-click the Oracle service.
 - b. Click the **General** tab.
 - c. Select **Manual** from the **Startup type** drop-down list and click **OK**.
3. Install Storage Mirroring on the target using the installation defaults. See the *HP OpenView Storage Mirroring getting started guide* for details.

Configure Storage Mirroring and begin mirroring and replication

1. Select **Start, Programs, Storage Mirroring, Management Console**.
2. Double-click your source machine to log on.
3. Right-click your source machine and select **New, Replication Set** and enter the desired name for the replication set.
4. Select the Oracle data you wish to protect. The table below will help you verify that you have selected all pertinent Oracle files necessary to enable high availability on your target machine.
 - **File Name**—This column lists each file name you need to select for replication. These files are assigned names during your initial Oracle installation and configuration causing some of these files to be unique to your environment. For example, if you used the default Oracle installation, your instance parameter file will be `init.ora` and located in the `\oracle\admin\instance_name\pfile` folder in the destination drive. If you supplied an instance name of `prod`, your `init.ora` file will be located in the `\oracle\admin\prod\pfile` folder by default. Files with unique names will be identified in the table with a wildcard asterisk (*).

Additionally, some of the files are identified with a number sign (#) placeholder in the file name. This placeholder identifies redundant files and/or multiple files which may exist because of tablespace configurations. For example, you may have control01.dbf, control02.dbf, undotbs01.dbf, and undotbs02.dbf.

- **File Description**—This column gives a brief description of each file.
- **Default Location and Table to Query**—The **Default Location** column identifies the directory where a specific file is stored if you choose the default installation directory. If you did not use the default directory and do not know the location of a specific file, query the table identified in the **Table to Query** column to determine the location. For detailed information on querying tables, see your Oracle reference guide. If you do not know the location of files that are not associated with a table, locate those files using the Windows Find feature.

| File Name | File Description | Default Location | Table to Query |
|----------------|-------------------------------|--|----------------|
| init.ora.xxxxx | Initialization Parameter File | \oracle\product\10.1.0\admin\instance_name\pfile | N/A |
| control#.ctl | Control File | \oracle\oradata\instance_name | V\$CONTROLFILE |
| redo#.log | Online Redo Log | \oracle\oradata\instance_name | V\$LOGFILE |
| system*.dbf | System Tablespace | \oracle\oradata\instance_name | V_\$DATAFILE |
| undotbs*.dbf | Undo Tablespace | \oracle\oradata\instance_name | V_\$DATAFILE |
| sysaux*.dbf | System Auxiliary Tablespace | \oracle\oradata\instance_name | V_\$DATAFILE |
| users#.dbf | Users Tablespace | \oracle\oradata\instance_name | V_\$DATAFILE |
| example*.dbf | Example Tablespace | \oracle\oradata\instance_name | V_\$DATAFILE |
| temp*.dbf | Temporary Tablespace | \oracle\oradata\instance_name | V_\$DATAFILE |
| listener.ora | Listener File | \oracle\product\10.1.0\db_1\network\admin | N/A |
| tnsnames.ora | Service Configuration File | \oracle\product\10.1.0\db_1\network\admin | N/A |
| pwd*.ora | Password Information File | \oracle\product\10.1.0\db_1\database | N/A |
| spfile*.ora | Server Parameter File | \oracle\product\10.1.0\db_1\database\ | N/A |



NOTE: There is a unique numeric identifier at the end of the `init.ora` file (xxxxx) that uniquely identifies initialization files when there is more than one installed.

Besides the initialization parameter file (`init.ora`), which is stored in a text format, administrators can choose to maintain the initialization parameters in a binary server parameter file (`spfile.ora`). This is installed by default

The Listener service will not start on the target if `listener.ora` is replicated from the source, but the target is not in a failed over state. This file contains configuration information that is not valid on the target unless it is in a failed over state. If you need to start Oracle on the target server without failing over, do not include `listener.ora` in the replication set. After you have selected all of the Oracle data you wish to protect, right-click the replication set name and select **Save** to save the replication set.

-
5. Right-click on the source and select **Properties**.
 6. On the Source tab, enable **Block Checksum All Files on a Difference Mirror** and click **OK**.
 7. Drag and drop the replication set onto the target. The Connection Manager dialog box opens.
 8. The **Source Server**, **Target Server**, **Replication Set**, and **Route** fields will automatically be populated. If you have multiple IP addresses on your target, verify the **Route** field is set to the correct network path. (For detailed information on connecting a source and target, see the *HP OpenView Storage Mirroring user's guide*.)
 9. Select **One to One** to map the replication set data from the source to an identical volume/directory structure on the target.
 10. Click **Connect** to start the mirror and replication processes.



NOTE: If you start Oracle and mount the replicated databases on the target, or if the data on the target is otherwise modified, the data on the source and target will no longer match. If the updated data on the target is not needed, perform a full or difference with block checksum mirror from the source to the target. If the updated data on the target is needed, restore the data from the target to the source.

Configure failover and begin failure monitoring

1. If a failure occurs, you will want to have the Oracle services start on the target machine automatically. To do this, create a batch file called `postover.bat` using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

NOTE: After you modify the sample scripts, save them with a new name to remove the `sample_` prefix. Copy the scripts to the directory where Storage Mirroring is installed.

The sample batch files provided are only examples. Because no two environments or configurations are exactly the same, you **MUST** modify the sample scripts in order to make the solution work in your environment.

SAMPLE_POSTOVER.BAT

```
rem ## Sample Oracle post-failover script
rem ## This script makes the following two assumptions:
rem ## 1. You have created a directory c:\oracle_fo
rem ## 2. The dbx_start.sql file referenced in this script is located
rem ##    in the oracle failover directory c:\oracle_fo
rem ## Note: %%a is a constant recognized in the command line mode and should NOT
rem ## be substituted.

rem ## The next two lines establish the working directories. Specify the location
rem ## of your Oracle program files in the first line. Do not leave any spaces
rem ## around the equal (=) sign.
set ORACLE_HOME=C:\ORACLE\ORA90
set ORACLE_FO=C:\ORACLE_FO

rem ## The following line CREATES the Oracle instance(s).
rem ## Substitute your oracle instance names for instance_name1 and instance_name2. Multiple instances
rem ## are separated by a space.
for %%a IN (instance_name1 instance_name2) DO %ORACLE_HOME%\bin\oradim -new -sid %%a -srvc
oracleservice%%a -startmode auto -pfile %ORACLE_HOME%\admin\instance_name\init.ora

rem ## The following line STARTS the Oracle instance(s).
rem ## Substitute your oracle instance names for instance_name1 and instance_name2. Multiple instances
rem ## are separated by a space.
for %%a IN (instance_name1 instance_name2) DO NET START ORACLESERVICE%%a

rem ## The following line starts other Oracle services. The services specified
rem ## are started on an as-needed basis depending on the Oracle operations
rem ## you are running. The Oracle HTTP Server and the Oracle Agent are started
rem ## in the following example. Add additional services as needed separating
rem ## the service names by a space. The outcome is redirected to a log file in
rem ## the failover directory. Check this log file for errors.
for %%a IN (OracleOraHomeHTTPServer OracleOraHomeAgent) DO NET START %%a
>>%ORACLE_FO%\start_services.log

rem ## The following lines mount and open the databases. Note that the same
rem ## command is repeated for each database. Also note that separate files
rem ## (DB1_Start.sql and DB2_Start.sql) are called for the different databases
rem ## so that the correct path to the parameter file (init.ora) is used.
rem ## The outcome is redirected to a log file in the failover directory. Check
rem ## this log file for errors. Substitute your instance names for instance_name1
rem ## and instance_name2.
set ORACLE_SID=instance_name1
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA"
<%ORACLE_FO%\db1_start.sql>%ORACLE_FO%\start_%ORACLE_SID%.log
set ORACLE_SID=instance_name2
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA"
<%ORACLE_FO%\db2_start.sql>%ORACLE_FO%\start_%ORACLE_SID%.log

rem ## The following line starts the Oracle Listener service. The outcome is
rem ## redirected to the failover directory where a log file is created.
rem ## Check this log file for errors.
NET START OracleOraHomeTNSListener >>%ORACLE_FO%\start_services.log
```



NOTE: After failover is complete, users will need to log back into the Oracle database. You may want to devise an automatic process to notify the users of this information and include it at the end of the post-failover script.

The following example shows the DB_START.SQL script used in the post-failover script.

SAMPLE_DBx_START.SQL

```
# Sample db_start.sql script called during the Storage Mirroring post-failover sequence.
# The following line identifies the Oracle instance to start.
# Replace instance_name with the name of the instance.
startup pfile=c:\oracle\admin\instance_name\pfile\init.ora
exit
```

2. After a failure is resolved, you will be ready to bring your source back online. At this time, you will want to stop the Oracle services on the target automatically. To do this, create a batch file called preback.bat using the sample batch file below. Save the batch file to the same directory where your Storage Mirroring files are installed.

SAMPLE_PREBACK.BAT

```
rem ## Sample Oracle pre-failback script

rem ## This script makes the following two assumptions:
rem ## 1. You have created a directory c:\oracle_fo
rem ## 2. The db_stop.sql file referenced in this script is located
rem ## in the oracle failover directory c:\oracle_fo
rem ## Note: %%a is a constant recognized in the command line mode and should NOT
rem ## be substituted.

rem ## The next two lines establish the working directories. Specify the location
rem ## of your Oracle program files in the first line. Do not leave any spaces
rem ## around the equal (=) sign.
set ORACLE_HOME=C:\ORACLE\ORA90
set ORACLE_FO=C:\ORACLE_FO

rem ## The following lines stop the Oracle databases.
rem ## Substitute your oracle instance names for instance_name1 and
rem ## instance_name2.
rem ## If you only have one Oracle database, remark out the
rem ## third and fourth lines of this section.
set ORACLE_SID=instance_name1
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA" <%ORACLE_FO%\db_stop.sql
>%ORACLE_FO%\stop_%ORACLE_SID%.log
set ORACLE_SID=instance_name2
%ORACLE_HOME%\bin\sqlplus.exe "/ AS SYSDBA" <%ORACLE_FO%\db_stop.sql
>%ORACLE_FO%\stop_%ORACLE_SID%.log

rem ## The following lines delete each instance service.
rem ## Substitute the name of the Oracle instance to delete for instance_name1
rem ## and instance_name2. Multiple instances are separated by a space.
FOR %%a IN (instance_name1 instance_name2) DO %ORACLE_HOME%\bin\oradim -delete -sid %%a

rem ## The following lines stop other non-instance specific Oracle services.
rem ## The value in parenthesis are the services to stop.
FOR %%a IN (OracleOraHomeTNSListener OracleOraHomeHTTPServer OracleOraHomeAgent) DO net stop %%a
>>%ORACLE_FO%\stop_service.log
```

The following example illustrates the DB_STOP.SQL script used in the pre-failback script.

SAMPLE_DB_STOP.SQL

```
# Sample db_stop.sql script called during the Storage Mirroring pre-failback script.
# Shutdown the database immediately.
shutdown immediate
exit
```

3. Select **Start, Programs, Storage Mirroring, Failover Control Center**.
4. Select the target machine from the list of available machines. If the target you need is not displayed, click **Add Target**, enter the machine name, and click **OK**.
5. To add a monitor for the selected target, click **Add Monitor**. Type the name of the source machine and click **OK**. The Monitor Settings window will open.
6. In the Monitor Settings window, mark the IP address that is going to failover.

7. Click **Scripts** and specify the scripts that were created previously using `postover.bat` for the target post-failover script and `preback.bat` for the target pre-failback script.
8. Click **OK** to go back to the Monitor Settings dialog box.
9. Click **OK** to begin monitoring the source machine.

In the event of a source machine failure, your target machine is now ready to stand in for your source. For details on monitoring failover, see the *HP OpenView Storage Mirroring user's guide*.

NOTE: In the event of a source machine failure and the target machine stands in, users that were connected during the time of the failure will be required to relogin to the database.

Restoring your Oracle data

If your source experiences a failure, such as a power, network, or disk failure, your target machine will stand in for the source while you resolve the source machine issues. During the source machine downtime, data is updated on the target machine. When your source machine is ready to come back online, the data is no longer current and must be updated with the new data on the target machine.

1. Verify that your source machine is not connected to the network. If it is, disconnect it.
2. Resolve the source machine problem that caused the failure.



NOTE: If you must rebuild your hard drive, continue with step 3. If you do not need to rebuild your hard drive, continue with step 6.

3. Install Windows. Since your source machine is not connected to the network, go ahead and use the source's original name and IP address.
4. Install Storage Mirroring using the installation defaults. .
5. Install Oracle using the installation settings previously used.
6. On the source, stop the Oracle service so that all Oracle files are closed on the source and the Storage Mirroring target can restore the changes back to the source.
7. Select **Start, Programs, Storage Mirroring, Failover Control Center**.
8. Select the target machine that is currently standing in for the failed source.
9. Select the failed source and click **Failback**. The pre-failback script entered during the failover configuration stops the Oracle services on the target so that no additional changes can be made.
10. You will be prompted to determine if you want to continue monitoring the source server. Do not choose **Continue** or **Stop** at this time.
11. Connect the source machine to the network.
12. After the source is back online, select whether or not you want to continue monitoring this source machine (**Continue** or **Stop**).



NOTE: Verify that the Storage Mirroring connection on the source has been disconnected (right-click the connection in the Storage Mirroring Management Console and select **Disconnect**).

13. To begin the restoration process, open the Storage Mirroring Management Console and select **Tools, Restoration Manager**.



NOTE: You can also run the Storage Mirroring DTCL automated restoration script, which can be found in the *HP OpenView Storage Mirroring user's guide*, to complete the remaining steps in this section.

14. Complete the appropriate fields as described below.

- **Original Source**—The name of the source machine where the data originally resided.
- **Restore From**—The name of the target machine that contains the replicated data.
- **Replication Set**—The name of the replication set to be restored.
- **Restore To**—The name of the machine where the data will be restored. This may or may not be the same as the original source machine.

15. Identify the correct drive mappings for the data and any other restoration options necessary. For detailed information on the restoration options, see the *HP OpenView Storage Mirroring user's guide*.

16. Verify that the selections you have made are correct and click **Restore**. The restoration procedure time will vary depending on the amount of data that you have to restore.

17. After the restoration is complete, start the Oracle services on the source machine.

18. Reestablish the Storage Mirroring Oracle replication set connection.

At this time, your data is restored back to your source machine, the source machine is again the primary Oracle server, and, if you selected to continue failover monitoring, the target is available to stand in for the source in the event of a failure.